

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Ranking Parser for a
Natural Language Processing System**

Inventor:
David Weise

ATTORNEY'S DOCKET NO. MS1-543US

1 **TECHNICAL FIELD**

2 This invention relates to ranking parses produced by a parser for a natural
3 language processing system.
4

5 **BACKGROUND**

6 In general, a computer is a digital machine that uses precise languages with
7 absolute values, such as "on", "off", "1", "0", "3 + 4", "AND", and "XOR". In
8 contrast, a human is an analog, biological machine that inherently uses imprecise
9 languages with few or no absolute values. Since computers are tools for human
10 use, input devices and input processing system are needed for humans to use the
11 computer tools.

12 Since it is generally easier to train humans to conform to the digital
13 requirements of computers than vice versa, humans have used precise input
14 interfaces such as a keyboard and a mouse. In addition, the computer is often only
15 required to receive the input and not to process it for syntax and semantics.

16 In the past, this has been the situation because of limited processing
17 capabilities of typical computers and because of the inherent difficulties of
18 modeling imprecise human language within a digital computer. However, as
19 typical computing power increases, natural language processing systems are being
20 used by computers to "understand" imprecise human language.
21

22 **Natural Language Processing**

23 A natural language processing (NLP) system is typically a computer-
24 implemented software system, which intelligently derives meaning and context
25

1 from an input string of natural language text. "Natural languages" are the
2 imprecise languages that are spoken by humans (e.g., English, French, Japanese).
3 Without specialized assistance, computers cannot distinguish linguistic
4 characteristics of natural language text. For instance, a sentence in a natural
5 language text read as follows:

6
7 Betty saw a bird.

8
9 A student of English understands that, within the context of this sentence,
10 the word "Betty" is a noun, the word "saw" is a verb, the word "a" is an adjective,
11 and the word "bird" is a noun. However, in the context of other sentences, the
12 same words might assume different parts of speech. Consider the following
13 sentence:

14
15 Use a saw.

16
17 The English student recognizes that the word "use" is a verb, the word "a"
18 is an adjective, and the word "saw" is a noun. Notice that the word "saw" is used
19 in the two sentences as different parts of speech—a verb and a noun—which an
20 English speaking person realizes. To a computer, however, the word "saw" is
21 represented by the same bit stream and hence can be identical for both sentences.
22 The computer is equally likely to consider the word "saw" as a noun as it is a verb,
23 in either sentence.
24
25

1 A NLP system assists the computer in distinguishing how words are used in
2 different contexts and in applying rules to construct intelligible language. A NLP
3 system has many different applications where a computer derives meaning and
4 information from the natural language of a human. Such applications include
5 speech recognition, handwriting recognition, grammar checking, spell checking,
6 formulating database searches, and language translation.

7 The core of a NLP system is its parser. Generally, a parser breaks an
8 utterance (such as a phrase or sentence) down into its component parts with an
9 explanation of the form, function, and syntactical relationship of each part.

10 11 NLP Parser

12 The NLP parser takes a phrase and builds for the computer a representation
13 of the syntax of the phrase that the computer can understand. A parser may
14 produce multiple different representations for a given phrase. The representation
15 makes explicit the role each word plays and the relationships between the words,
16 much in the same way as grade school children diagram sentences. In addition to
17 "diagramming" a sentence, the parser ranks the multiple diagrams in order of most
18 likely meaning to least likely.

19 Herein, an utterance is equivalent to a phrase. A phrase is a sequence of
20 words intended to have meaning. In addition, a sentence is understood to be one or
21 more phrases. In addition, references herein to a human speaker include a writer
22 and speech includes writing.
23
24
25

1 Fig. 1 shows a NLP parser 20 of a typical NLP system. The parser 20 has
2 four key components:

- 3 • Tokenizer 28;
- 4 • Grammar Rules Interpreter 26;
- 5 • Searcher 30; and
- 6 • Parse Ranker 34.

7 The parser 20 receives a textual string 22. Typically, this is a sentence or a
8 phrase. The parser also receives grammar rules 24. These rules attempt to codify
9 and interpret the actual grammar rules of a particular natural language, such as
10 English. Alternatively, these rules may be stored in memory within the parser.

11 The grammar rules interpreter 26 interprets the codified grammar rules. The
12 tokenizer 28 identifies the words in the textual string 22, looks them up in a
13 dictionary, makes records for the parts of speech (POS) of a word, and passes
14 these to the searcher.

15 The searcher 30 in cooperation with the grammar rules interpreter generates
16 multiple grammatically correct parses of the textual string. The searcher sends its
17 results to the parse ranker 34.

18 The parse ranker 34 mathematically measures the “goodness” of each parse
19 and ranks them. “Goodness” is a measure of the likelihood that such a parse
20 represents the intended meaning of the human speaker (or writer). The ranked
21 output of the parser ranker is the output of the ranker. This output is one or more
22 of parses 38 ranked from most to least goodness.
23
24
25

Foundational Concepts

Three concepts form the foundation for understanding the invention described herein: statistics, linguistics, and computational linguistics.

Statistics is the branch of mathematics that deals with the relationships among and between groups of measurements, and with the relevance of similarities and differences in those relationships.

Linguistics is the analytic study of human natural language.

Computational linguistics is the analytic study of human natural language within computer science to mathematically represent language rules such as grammar, syntax, and semantics.

Statistics

Probability. The expression "Prob(x)" is the probability of event x occurring. The result of Prob(x) is a number between zero (0) and one (1), where zero means that the event never occurs and one means that it always occurs. For example, using a six-sided fair die with the sides labeled 1-6, the probability of rolling a three is 1/6. Similarly, using a randomly shuffled deck of cards (All examples using a deck of cards are based upon a standard American deck of cards having four suits (spades, hearts, diamonds, clubs) and thirteen cards per suit:

$$\text{Prob}(\text{top card is an Ace}) = 1/13$$

$$\text{Prob}(\text{top card is a club}) = 1/4$$

$$\text{Prob}(\text{top card is 3 of diamonds}) = 1/52$$

Estimating Probabilities using Training Data. The probability of events using a randomly shuffled deck or fair die can be mathematically derived. However, in many cases, there is no mathematical formula for a probability of a given event. For example, assume that one wished to determine the probability of rolling a three given a weighted die. The probability may be $1/6$ (as it would be with a fair die), but it is likely to be more or less than that depending upon how the die is weighted.

How would one estimate the probability? The answer is to run an experiment. The die is thrown many times and the number of rolls where “3” is rolled is counted. This data is called the “training data”. It is sometimes called the “training corpus.” To determine the probability of rolling a three in the future, it is assumed that the behavior of the die in the future will be the same as it was during our experiment and thus:

$$\text{Prob(event)} = \frac{\text{Count(event)}}{\text{Total number of events}}$$

$$\text{Prob(roll is a 3)} = \frac{\text{Count(number of 3 rolls in the experiment)}}{\text{Total \# of rolls in the experiment}}$$

1 In general, the accuracy of the estimate increases as the amount of training
2 data increases. Theoretically, the estimates increase in accuracy as the amount of
3 training data increases.

4
5 Conditional Probability. Conditional probabilities are used when there is
6 additional information known about an event that affects the likelihood of the
7 outcome. The notation used is $\text{Prob}(x | y)$ meaning, "What is the probability of an
8 unknown event x occurring given that known event y occurred."

9 Conditional probability is defined to be:

$$\begin{aligned} \text{Prob}(x|y) &= \frac{\text{Prob}(x\&y)}{\text{Prob}(y)} \\ &= \frac{\text{Count}(x\&y)}{\text{Total \# of events}} \\ &= \frac{\text{Count}(y)}{\text{Total \# of events}} \\ &= \frac{\text{Count}(x\&y)}{\text{Count}(y)} \end{aligned}$$

22 When the known event is predicative of the outcome, knowing the
23 conditional probabilities is better than knowing just the unconditional probability.
24 For example, assume that a man is in a casino playing the following game. The
25

1 man then can bet \$1 or pass. The House rolls a pair of dice. If the man bets and if
2 the dice sum to 12, the man gets \$35, otherwise the man loses his bet. Since the
3 probability of rolling two die that sum to 12 is 1/36, the man should expect to lose
4 money playing this game. On average, the man will make only \$35 for every \$36
5 that he bets.

6 Now suppose that the man had a fairy godmother that could whisper in his
7 ear and tell him whether one of the die rolled was going to be a six. Knowing this,
8 the probabilities of rolling a twelve are:

9
10
$$\text{Prob}(\text{two die summing to 12} \mid \text{one die is a 6}) = 1/6$$

11
$$\text{Prob}(\text{two die summing to 12} \mid \text{neither die is a 6}) = 0$$

12

13 With the fairy godmother's help, the man can make money on the game.
14 The strategy is to only bet when the fairy godmother says that one of the die is a
15 six. On average, the man should expect to make \$35 for every \$6 that he bets.

16 As another example, consider the problem of predicting what the next word
17 in a stream of text will be. E.g., is the next word after "home" more likely to be
18 "table" or "run."? Word_i represents the i^{th} word in the lexicon and $\text{Prob}(\text{word}_i)$ is
19 the probability that word_i will be the next word in the stream. The standard
20 approach (using unconditional probability) for computing $\text{Prob}(\text{word}_i)$ is to take a
21 training corpus and count up the number of times the word appears. This formula
22 represents this approach:

23
24
$$\text{Prob}(\text{word}_i) = \frac{\text{Count}(\text{word}_i)}{\quad}$$

25

Better results are achieved by using conditional probabilities. In English, words don't appear in a random order. For example, "Table red the is" is highly unlikely but "The table is red" is common. Put another way, the last word in a stream is predicative of the next word that will follow. For example, if the last word in a stream is "the" or "a", then the next word is usually either an adjective or noun, but is rarely a verb. If the last word in a stream is "clever", then the next word is likely to be an animate noun like "boy" or "dog" and not likely to be an inanimate noun like "stone."

By making use of this information, the next word that will appear in a stream of text may be better predicted. If every pair of words in our lexicon is considered and a record of how often that pairs appear is kept, the probability of a specific word pairing is:

$$\text{Prob}(\text{word}_i \mid \text{word}_k \text{ appears before word}_i) = \frac{\text{Count}(\text{word}_i \text{ and } (\text{word}_k \text{ before word}_i))}{\text{Count}(\text{word}_k)}$$

This conditional probability is much more accurate than $\text{Prob}(\text{word}_i)$. This technique is commonly used by conventionally speech recognizers to predict what words are likely to follow a given speech fragment.

Sparse Data Problem. The sparse data problem occurs when there is not enough data in a training corpus to distinguish between events that never occur

1 versus events that are possible but just didn't happen to occur in the training
2 corpus. For example, $\text{Prob}(\text{word}_i + 1 \mid \text{word}_i)$ is being computed by counting how
3 often pairs of words occur in the corpus. If in the training corpus the pair of words
4 "gigantic" and "car" never appears, it would be wrong to conclude that it is
5 impossible in the English language to have the words "gigantic" and "car"
6 together.

7 In general, natural languages have nearly an infinite number of possible
8 word, phrase, and sentence combinations. The training corpus used to determine
9 conditional probabilities must necessarily be a subset of this set of infinite
10 combinations. Thus, the sparse data problem results in poor probabilities with a
11 given combination when the training corpus did not include that given
12 combination.

13 Chain Rule.

$$14 \quad \text{Prob}(a, b \mid c) = \text{Prob}(a \mid c) \text{Prob}(b \mid a, c)$$

15 Linguistics

16 Linguistics is the scientific study of language. It endeavors to answer the
17 question--what is language and how it is represented in the mind? Linguistics
18 focuses on describing and explaining language.

19 Linguistics focuses on languages' syntax (sentence and phrase structures),
20 morphology (word formation), and semantics (meaning). Before a computer
21
22
23
24
25

1 representation model of a natural language can be generated and effectively used,
2 the natural language must be analyzed. This is the role of linguistics.

3
4 Part of Speech. Linguists group words of a language into classes, which
5 show similar syntactic behavior, and often a typical semantic type. These word
6 classes are otherwise called “syntactic” or “grammatical categories”, but more
7 commonly still by the traditional names “part of speech” (POS). For example,
8 common POS categories for English include noun, verb, adjective, preposition,
9 and adverb

10
11 Word Order and Phrases. Words do not occur in just any order. Languages
12 have constraints on the word order. Generally, words are organized into phrases,
13 which are groupings of words that are clumped as a unit. Syntax is the study of
14 the regularities and constraints of word order and phrase structure. Among the
15 major phrase types are noun phrases, verb phrases, prepositional phrases, and
16 adjective phrases.

17
18 Headword. The headword is the key word in a phrase. This is because it
19 determines the syntactic character of a phrase. In a noun phrase, the headword is
20 the noun. In a verb phrase, it is the main verb. For example, in the noun phrase
21 “red book”, the headword is “book.” Similarly, for the verb phrase “going to the
22 big store”, the headword is “going.”

1 Modifying Headword. A modifying headword is the headword of a sub-
2 phrase within a phrase where the sub-phrase modifies the main headword of the
3 main phrase. Assume a phrase (P) has a headword (hwP) and a modifying sub-
4 phrase (M) within the P that modifies hwP. The modifying headword (hwM) is
5 the headword of this modify phrase (M).

6 For example, if the phrase is “The red bear growled at me”, the headword is
7 “growled,” the modifying phrase is “the red bear,” and the modifying headword is
8 “bear.” If the phrase is “running to the store”, then the headword is “running”, the
9 modifying phrase is “to the store”, and the modifying headword is “to.”

10
11 Syntactic Features. Syntactic features are distinctive properties of a word
12 relating to how the word is used syntactically. For example, the syntactic features
13 of a noun include whether it is singular (e.g. cat) or plural (e.g. cats) and whether it
14 is countable (e.g. five forks) or uncountable (e.g. air). The syntactic feature of a
15 verb includes whether or not it takes an object:

- 16 • Intransitive verbs do not take an object. For example, “John
17 laughed,” and “Bill walked,”
- 18 • Mono-transitive verbs take a single direct object. For example, “I hit
19 the ball”,
- 20 • Di-transitive verbs takes a direct and an indirect object. For example,
21 “I gave Bill the ball,” and “I promised Bill the money”

Computational Linguistics

Transitions (i.e., Rewrite Rules). The regularities of a natural language's word order and grammar are often captured by a set of rules called "transitions" or "rewrite rules." The rewrite rules are a computer representation of rules of grammar. These transitions are used to parse a phrase.

A rewrite rule has the notation form: "symbolA \rightarrow symbolB symbolC ...". This indicates that symbol (symbolA) on the left side of the rule may be rewritten as one or more symbols (symbolB, symbolC, etc.) on the right side of the rule.

For example, symbolA may be "s" to indicate the "start" of the sentence analysis. SymbolB may be "np" for noun phrase and symbolC may be "vp" for verb phrase. The "np" and "vp" symbols may be further broken down until the actual words in the sentence are represented by symbolB, symbolC, etc.

For convenience, transitions can be named so that the entire rule need not be recited each time a particular transition is referenced. In Table 1 below the name of the transitions are provided under the "Name" heading. The actual transitions are provided under the "Transition" heading. Table 1 provides an example of transitions being used to parse a sentence like "Swat flies like ants":

Name	Transition (i.e., rewrite rule)
s_npvp	s → np vp
s_vp	s → vp
np_noun	np → noun
np_nounpp	np → noun pp
np_nounnp	np → noun np
vp_verb	vp → verb
vp_verbnp	vp → verb np
vp_verbpp	vp → verb pp
vp_verbnppp	vp → verb np pp
pp_prepnp	pp → prep np
prep_like	prep → like
verb_swat	verb → swat
verb_flies	verb → flies
verb_like	verb → like
noun_swat	noun → swat
noun_flies	noun → flies
noun_ants	noun → ants

Table 1

In Table 1 above, the transition names (on the left-hand column) represent and identify the transition rule (on the right-hand column). For example, “np_nounpp” is the name for “np → noun pp” rule, which means that “noun phrase” may be rewritten as “noun” and “prepositional phrase.”

1 Context Free Grammar (CFG). The nature of the rewrite rules is that a
2 certain syntactic category (e.g, noun, np, vp, pp) can be rewritten as one or more
3 other syntactic categories or words. The possibilities for rewriting depend solely
4 on the category, and not on any surrounding context, so such phrase structure
5 grammars are commonly referred to as context-free grammars (CFG).

6 Fig. 2 illustrates a CFG parse tree 50 of a phrase (or sentence). This tree-
7 like representation of the sentence “flies like ants” is deconstructed using a CFG
8 set of rewrite rules (i.e, transitions). The tree 50 has leaf nodes (such as 52a-52c
9 and 54a-54g.)

10 The tree 50 includes a set of terminal nodes 52a-52c. These nodes are at
11 the end of each branch of the tree and cannot be further expanded. For example,
12 “like” 52b cannot be expanded any further because it is the word itself.

13 The tree 50 also includes a set of non-terminal nodes 54a-54g. These nodes
14 are internal and may be further expanded. Each non-terminal node has immediate
15 children, which form a branch (i.e., “local tree”). Each branch corresponds to the
16 application of a transition. For example, “np” 54b can be further expanded into a
17 “noun” by application of the “np_noun” transition.

18 Each non-terminal node in the parse tree is created via the application of
19 some rewrite rule. For example, in Fig. 2, the root node 54a was created by the
20 “s→np vp” rule. The “VP” node 54d by the “s→verb np” rule.

21 The tree 50 has a non-terminal node 54a designated as the starting node and
22 it is labeled “s.”
23
24
25

1 In general, the order of the children in each branch generates the word order
2 of the sentence, and the tree has a single root node (in Fig. 2 it is node 54a), which
3 is the start of the parse tree.
4

5 Segtypes. A non-terminal node has a type that is called its "segtype." In
6 Fig. 2, each non-terminal node 54a-g is labeled with its segtype. A node's segtype
7 identifies the rule that was used to create the node (working up from the terminal
8 nodes). In Table 1 above, the segtypes are shown under the "Transition" heading
9 and to the left of the "→" symbol. For example, the segtype of node 54b in Fig. 2
10 is "np" because the rule "np → noun" was used to create the node.

11 In given grammar, a segtype can be many different values including, for
12 example: NOUN, NP (noun phrase), VERB, VP (verb phrase), ADJ (adjective),
13 ADJP (adjective phrase), ADV (adverb), PREP (preposition), PP (prepositional
14 phrase), INFCL (infinitive clauses), PRPRT (present participial clause) PTPRT
15 (past participial clause), RELCL (relative clauses), and AVPVP (a verb phrase that
16 has a verb phrase as its head).
17

18 Node-Associated Functional Notation. In this document, a functional
19 notation is used to refer to the information associated with a node. For example, if
20 a variable "n" represents a node in the tree, then "hw(n)" is the headword of node
21 "n."
22
23
24
25

1 The following functions are used through out this document:

- 2 • hw(n) is the headword of node n
- 3 • segtype(n) is the segtype of node n
- 4 • trans(n) is the transition (rewrite rule) associated with node n (e.g.,
- 5 the rules under the heading “Transition” in Table 1)
- 6 • trn(n) is the name of the transition (e.g. the names under the heading
- 7 “Name” in Table 1)
- 8 • modhw(n) is the modifying headword of node n

9
10 Annotated Parse Tree. A parse tree can be annotated with information
11 computed during the parsing process. A common form of this is the lexicalized
12 parse tree where each node is annotated with its headword. One can annotate a
13 parse tree with additional linguistic information (e.g. syntactic features).

14 Fig. 3 shows an example of such a lexicalized parse tree 60. (For the
15 purposes of this example, directional path 66 with circled reference points is
16 ignored.) Fig. 3 is a parse tree of one or many parses of the sentence, “swat flies
17 like ants.” Terminal nodes 62a-d, which are the words of the sentence, are not
18 annotated. Non-terminal nodes 64a-i are annotated. For example, node 64h has a
19 segtype of “noun” and is annotated with “hw=ants”. This means that its headword
20 is “ants.”

21 The parse tree 60 in Fig. 3 is also annotated with the names of the
22 transitions between nodes. For example, the transition name “vp_verbvp” is listed
23 between node 64f and node 64h.
24
25

1 Probabilistic Context Free Grammar (PCFG). A PCFG is a context free
2 grammar where every transition is assigned a probability from zero to one.
3 PCFGs have commonly been used to define a parser's "goodness" function.
4 "Goodness" is a calculated measurement of the likelihood that a parse represents
5 the intended meaning of the human speaker. In a PCFG, trees containing
6 transitions that are more probable are preferred over trees that contain less
7 probable transitions.

8 Since the probability of a transition occurring cannot be mathematically
9 derived, the standard approach is to estimate the probabilities based upon a
10 training corpus. A training corpus is a body of sentences and phrases that are
11 intended to represent "typical" human speech in a natural language. The speech
12 may be intended to be "typical" for general applications, specific applications,
13 and/or customized applications. This "training corpus" may also be called
14 "training data."

15 Thus, the probabilities are empirically derived from analyzing a training
16 corpus. Various approaches exist for doing this. One of the simplest approaches is
17 to use an unconditional probability formula like this:

$$\begin{array}{lcl} \text{Prob(event) =} & \frac{\text{Count(event)}}{\text{Total number of events}} \\ \\ \text{Prob(trans}_i\text{) =} & \frac{\text{Count(times trans}_i\text{ appears)}}{\text{Total \# of transitions in the training corpus}} \end{array}$$

1 However, this approach, by itself, produces inaccurate results because the
2 likelihood that a transition will apply is highly dependent upon the current
3 linguistic context, but this approach does not consider the current linguistic
4 context. This approach simply considers occurrences of specific transitions
5 (trans_i).
6

7 Depth First Tree Walk. In order to analyze each node of a parse tree to
8 rank parse trees, a parser must have a method of “visiting” each node. In other
9 words, the nodes are examined in a particular order.
10

11 A “depth first tree walk” is a typical method of visiting all the nodes in a
12 parse tree. In such a walk, all of a node’s children are visited before any of the
13 node’s siblings. The visitation is typically from top of the tree (i.e., the start node)
14 to the bottom of the tree (i.e., terminal nodes). Such visitation is typically done
15 from left-to-right to correspond to the order of reading/writing in English, but may
16 be done from right-to-left.

17 The directional path 66 of Fig. 3 shows a depth first tree walk of the parse
18 tree. The sequence of the walk is shown by the directional path 66 with circled
19 reference points. The order of the stops along the path is numbered from 1 to 13
20 by the circled reference points.

21 Generative model of syntax. Each sentence-tree pair in a language has an
22 associated top-down derivation consisting of a sequence of rule applications
23 (transitions) of a grammar.
24
25

1 Augmented Phrase Structured Grammar (APSG). An APSG is a CFG that
2 gives multiple names to each rule, thereby limiting the application of each
3 “named” rule. Thus, for each given rewrite rule there are more than one name and
4 the name limits its use to specific and narrower situations. For example, the
5 structure “VP → NP VP” may have these limiting labels: “SubjPQuant” and
6 “VPwNPI.”

7 SubjPQuant specifies subject post-quantifiers on a verb phrase. For
8 example, in “all found useful...”, “all” is a subject post-quantifier, In “we all
9 found useful the guidelines” is [NP all][VP found useful the guidelines].
10 VPwNPI specifies a subject to a verb phrase. For example, in “John hit the ball”
11 [NP John] [VP hit the ball] where John is the subject.

12 The Problem

13 Given the ambiguity that exists in natural languages, many sentences have
14 multiple syntactic interpretations. The different syntactic interpretations generally
15 have different semantic interpretations. In other words, a sentence has more than
16 one grammatically valid structure (“syntactic interpretation”) and as a result, may
17 have more than one reasonable meaning (“semantic interpretation”). A classic
18 example of this is the sentence, “time flies like an arrow.” There are seven valid
19 syntactic parse trees.
20

21 Figs. 4a and 4b show examples of two of the seven valid parses of this
22 sentence. For the parse tree 70 of Fig. 4a, the object “time” 74 moves in a way
23 that is similar to an arrow. For the parse tree 80 of Fig. 4b, the insects called
24
25

1 “time flies” 84 enjoy the arrow object; just as one would say “Fruit flies like a
2 meal.”

3 Either parse could be what the speaker intended. In addition, five other
4 syntactically valid parses may represent the meaning that the speaker intended.

5 How does a NLP system determine which parse is the “correct” one. It is
6 better to say the most “correct” one. How does a NLP parser judge amongst the
7 multiple grammatically valid parses and select the most “correct” parse?

8 9 **Previous Approaches**

10 Generally. A parser needs a way to accurately and efficiently rank these
11 parse trees. In other words, the parser needs to compute which parse tree is the
12 most likely interpretation for a sentence, such as “time flies like an arrow.”

13 Since human language is inherently imprecise, rarely is one parse one
14 hundred percent (100%) correct and the others never correct. Therefore, a parser
15 typically ranks the parses from most likely to be correct to least likely to be
16 correct. Correctness in this situation is a measure of what a human most likely
17 means by a particular utterance.

18 A conventional approach is to use a “goodness” function to calculate a
19 “goodness measure” of each valid parse. Existing parsers differ in the extent to
20 which they rely on a goodness function, but most parsers utilize one.

21 A simple parser may generate all possible trees without regard to any
22 linguistic knowledge and then allow the goodness function to do all the work in
23 selecting the desired parse. Alternatively, a parser generates reasonable trees
24 based on linguistic knowledge and then uses the goodness function to choose
25

1 between the reasonable trees. In either case, the problem is to implement an
2 efficient goodness function that accurately reflects and measures the most likely
3 meaning of an utterance.

4
5 Straw Man Approach. The most straightforward approach is the “straw
6 man approach.” The goodness function of this approach computes the probability
7 of a given parse tree based upon how often identical trees appeared in a training
8 corpus. This approach is theoretical and is rarely (if ever) used in practice. This is
9 because it is inaccurate without an impractically huge training corpus that
10 accurately represents nearly all-possible syntactic and semantic constructions
11 within a given language.

12 Using the straw man approach, the probability of a parse tree is defined to
13 be:

$$\text{Prob}(\text{parse}) = \frac{\text{Count}(\text{parse})}{\text{Total \# of trees in the training corpus}}$$

14
15
16
17
18 For example, assume in the training corpus the sentence, “time flies like an
19 arrow” appears ten times. The parse represented by the parse tree 70 of Fig. 4a
20 appears in nine of those times. In addition, the parse represented by the parse tree
21 80 of Fig. 4b appears only once. Thus, the probability of the parse tree 70 of Fig.
22 4a would be ninety percent.

1 If the parse of parse tree 70 is the correct parse, then this example provide
2 good results. Note that, the exact sentence had to occur multiple times within the
3 corpus to provide such good results.

4 Theoretically, given enough training data, the straw man approach can be
5 highly accurate. However, the amount of training data required is astronomical.
6 First, it requires that the tagged training corpus contain all the sentences that the
7 parser is likely to ever encounter. Second, the sentences must appear in the correct
8 ratios corresponding to their appearance within the normal usage of the natural
9 language. In other words, common sentences occurring more often than
10 uncommon sentences and in the right proportion.

11 Creating such a huge training corpus is infeasible. However, working from
12 a smaller corpus creates sparse data problems.

13
14 Statistical Hodgepodge Approach. Using this approach, the goodness of a
15 parse may be determined by a collection of mostly unrelated statistical
16 calculations based upon parts of speech, syntactic features, word probabilities, and
17 selected heuristic rules.

18 A goodness function using such an approach is utilized by the grammar
19 checker in "Office 97" by the Microsoft Corporation. Parses were assigned a
20 score based upon statistical information and heuristic rules. These scores were
21 often called "POD" scores.

22 Since this hodgepodge approach employs heuristics and does not use a
23 unifying methodology for calculating the goodness measure of parses, there are
24 unpredictable and unanticipated results that incorrectly rank the parses.
25

1
2 Syntactic Bigrams Approach. This approach uses collocations to compute
3 a goodness function. A collocation is two or more words in some adjacent
4 ordering or syntactic relationship. Examples of such include: “strong tea”,
5 “weapons of mass destruction”, “make up”, “the rich and powerful”, “stiff
6 breeze”, and “broad daylight.”

7 Specifically, syntactic bigrams are two-word collocation. The basic idea is
8 to find the probability of two words of being in a syntactic relationship to each
9 other, regardless of where those words appear in the sentence. The words may be
10 adjacent (e.g., “I drink coffee.”), but need not be (e.g., “I love to drink hot black
11 coffee.”) For example, the object of the verb “drink” is more likely to be “coffee”
12 or “beer” than “table”. This can be used to create a goodness function based on
13 “syntactic bigrams.”

14 If the following four sentences appeared in the training corpus, all four
15 would provide evidence that “coffee” is often the object of the verb “to drink”:
16

17 I drink coffee.

18 I drink black coffee.

19 I love to drink hot black coffee.

20 I drink, on most days of the week, coffee in the morning.
21
22

23 However, because of the huge potential number of word combinations, this
24 approach requires a hefty training corpus.
25

Transition Probability Approach (TPA). A goodness function may be calculated using a generative grammar approach. Each sentence has a top-down derivation consisting of a sequence of rule applications (transitions). The probability of the parse tree is defined to be the product of the probabilities of the transitions.

There are a number of different ways to assign probabilities to the transitions. For this example, the transition probabilities are conditioned on segtype:

Prob(parse)

$$= \prod_i \text{Prob}(n_i)$$

$$= \prod_i \text{Prob}(\text{trans}(n_i) | \text{segtype}(n_i))$$

$$= \prod_i \frac{\text{Count}(\text{trans}(n_i) \ \& \ \text{segtype}(n_i))}{\text{Count}(\text{segtype}(n_i))}$$

Where

n_i : is the i^{th} node

$\text{trans}(n_i)$: is the transition out of n_i of the form $X \rightarrow Y Z$

$\text{segtype}(n_i)$: is the segtype of n_i

\prod_i is the notation to combine (e.g., multiply) over all nodes i in the parse

tree

1
2 For example, suppose that probabilities are assigned to each transition
3 shown in Table 1 above and those probabilities are based upon some training
4 corpus. The training corpus would contain parsed sentences such that the system
5 can count the number of times each transition occurred. In other words, the
6 system counts the number of times each particular grammar rule was used to
7 generate the parse. The result might be:
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

Transition	Count	Prob(trans segtype)	
s → np vp	80	.8	Sum = 1.0
s → vp	20	.2	
np → noun	80	.4	Sum = 1.0
np → noun pp	100	.5	
np → noun np	20	.1	
vp → verb	40	.4	Sum = 1.0
vp → verb np	20	.2	
vp → verb pp	20	.2	
vp → verb np pp	20	.2	
pp → prep np	10	1	Sum = 1.0
prep → like	10	1	Sum = 1.0
verb → swat	10	.1	Sum = 1.0
verb → flies	50	.5	
verb → like	40	.4	
noun → swat	100	.5	Sum = 1.0
noun → flies	50	.25	
noun → ants	50	.25	

Table 2

Using the PCFG represented by Table 2 above, the probability of a parse tree can be computed below as follows:

$$\begin{aligned}\text{Prob}(S) &= \text{Prob}(s \rightarrow np \text{ } vp) * \\ &\quad \text{Prob}(np \rightarrow \text{noun } np) * \\ &\quad \text{Prob}(\text{noun} \rightarrow \text{wat}) * \\ &\quad \text{Prob}(np \rightarrow \text{noun}) * \\ &\quad \text{Prob}(\text{noun} \rightarrow \text{flies}) * \\ &\quad \text{Prob}(vp \rightarrow \text{verb } np) * \\ &\quad \text{Prob}(\text{verb} \rightarrow \text{like}) * \\ &\quad \text{Prob}(np \rightarrow \text{noun}) * \\ &\quad \text{Prob}(\text{noun} \rightarrow \text{ants}) \\ &= .8 * .4 * .05 * .4 * .4 * .45 * .3 * .4 * .4 * .5 \\ &= 0.000027648\end{aligned}$$

However, this approach does not define a very accurate goodness function. Alone, a PCFG is generally poor at ranking parses correctly. A PCFG prefers common constructions in a language over less common ones.

Ancestor Dependency-Based Generative Approach (ADBGA). This approach assumes a top-down, generative grammar approach. It defines a formulism for computing the probability of a transition given an arbitrary set of linguistic features. Features might include headword, segtype, and grammatical number, though the formulism is independent of the actual features used. This approach does not attempt to define a particular set of features.

1 A transition is assumed to have the form:

2
3
$$(a_1, a_2, \dots a_g) \rightarrow (b_1, b_2, \dots b_g) (c_1, c_2, \dots c_g)$$

4
5 where

6 $a_1, a_2, \dots a_g$ are the features of a parent node

7 $b_1, b_2, \dots b_g$ are the features of a left child

8 $c_1, c_2, \dots c_g$ are the features of a right child

9
10 The probability of a transition is

11
12
$$\text{Prob}(b_1, b_2, \dots b_g, c_1, c_2, \dots c_g \mid a_1, a_2, \dots a_g)$$

13
14 Using the chain rule, this approach then conditions each feature on the
15 parent feature and all features earlier in the sequence:

16
17
$$\text{Prob}(b_1, b_2, \dots b_g, c_1, c_2, \dots c_g \mid a_1, a_2, \dots a_g)$$

18
$$= \text{Prob}(b_1 \mid a_1, \dots, a_g) *$$

19
$$\text{Prob}(b_2 \mid a_1, \dots, a_g, b_1) *$$

20
$$\text{Prob}(b_3 \mid a_1, \dots, a_g, b_1, b_2) *$$

21
$$\dots *$$

22
$$\text{Prob}(b_g \mid a_1, \dots, a_g, b_1, \dots, b_{g-1}) *$$

23
$$\text{Prob}(c_1 \mid a_1, \dots, a_g, b_1, \dots, b_g) *$$

24
$$\text{Prob}(c_2 \mid a_1, \dots, a_g, b_1, \dots, b_g, c_1) *$$

*

...

$$\text{Prob}(c_g \mid a_1, \dots, a_g, b_1, \dots, b_g, c_1, \dots, c_{g-1})$$

Background Summary

It is desirable for a NLP parser to be able to computationally choose the most probable parse from the potentially large number of possible parses. For example, the sentence "Feeding deer prohibited" may be logically interpreted to mean either the act of feeding is prohibited or that a type deer is prohibited.

A parser typically uses a goodness function to generate a "goodness measure" that ranks the parse trees. Conventional implementations use heuristic ("rule of thumb") rules and/or statistics based on the part of speech of the words in the sentence and immediate syntactic context.

The goodness function is a key component to a NLP parser. By improving the goodness function, the parser improves its accuracy. In particular, the goodness function enables the parser to choose the best parse for an utterance. Each parse may be viewed as a tree with branches that eventually branch to each word in a sentence.

Existing NLP parsers rank each parse tree using conventional goodness measures. To determine the parse with the highest probability of being correct (i.e., the highest goodness measure), each branch of each parse tree is given a probability. These probabilities are generated based upon a large database of correctly parsed sentences (i.e., "training corpus"). The goodness measure of each parse tree is then calculated by combining assigned probabilities of each branch in

1 branch in a parse tree. This conventional statistical goodness approach is typically
2 done with little or no consideration for contextual words and phrases.
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

SUMMARY

A natural language parse ranker of a natural language processing (NLP) system employs a goodness function to rank the possible grammatically valid parses of an utterance. The goodness function generates a statistical goodness measure (SGM) for each valid parse. The parse ranker orders the parses based upon their SGM values. It presents the parse with the greatest SGM value as the one that most likely represents the intended meaning of the speaker. The goodness function of this parse ranker is highly accurate in representing the intended meaning of a speaker. It also has reasonable training data requirements.

With this parse ranker, the SGM of a particular parse is the combination of all of the probabilities of each node within the parse tree of such parse. The probability at a given node is the probability of taking a transition ("grammar rule") at that point. The probability at a node is conditioned on highly predicative linguistic phenomena. Such phenomena include headwords, "phrase level", and "syntactic history." Transitions, headwords (including modifying headwords), phrase level", and syntactic history are collectively called "linguistic features."

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a schematic illustration of an exemplary natural language processing system.

Fig. 2 is an illustration of a typical parse tree representing a syntactically valid parse of sample phrase, "flies like ants."

1 Fig. 3 is another illustration of a typical parse tree representing a
2 syntactically valid parse of sample phrase, "swat flies like ants." This parse tree is
3 annotated to indicate transitions and headwords.

4 Figs. 4a and 4b illustrate two exemplary parse trees of two of seven
5 syntactically valid parses of sample phrase, "time flies like an arrow."

6 Figs. 5a and 5b show fragments of two pairs of typical parse trees. The
7 parse tree of Fig. 5a does not use headword annotation, but the parse tree of Fig.
8 5b does.

9 Figs. 6a and 6b show fragments of two pairs of typical parse trees. The
10 parse tree of Fig. 6a shows a parse done in accordance with an exemplary
11 grammar. The parse tree of Fig. 6b shows a parse tree that includes a null
12 transition.

13 Fig. 7 shows fragments of a pair of typical parse trees and illustrates the use
14 of syntactic history.

15 Fig. 8 shows a typical parse tree of a sample sentence, "Graceland, I like to
16 visit." This figure illustrates the "topicalization" syntactic phenomenon.

17 Fig. 9 shows a fragment of a genericized parse tree. This figure illustrates
18 what is known and not known at a node.

19 Fig. 10 is a flowchart illustrating the methodology of an implementation of
20 the training phase of the exemplary parser.

21 Fig. 11 is a flowchart illustrating the methodology of an implementation of
22 the run-time phase of the exemplary parser.

23 Fig. 12 is an example of a computing operating environment capable of
24 implementing the exemplary ranking parser for NLP.
25

DETAILED DESCRIPTION

The following description sets forth a specific embodiment of the ranking parser for natural language processing (NLP) that incorporates elements recited in the appended claims. This embodiment is described with specificity in order to meet statutory written description, enablement, and best-mode requirements. However, the description itself is not intended to limit the scope of this patent. Rather, the inventor has contemplated that the claimed ranking parser might also be embodied in other ways, in conjunction with other present or future technologies.

The exemplary ranking parser described herein may be implemented by a program submodule of a natural language processing (NLP) program module. It may also be implemented by a device within a NLP device. For example, a parse ranker 34 in Fig. 1 may be a program module implementing the exemplary parser within a NLP program system 20. Alternatively, the parse ranker 34 in Fig. 1 may be a parse ranker 34 in Fig. 1 may be a device implementing the exemplary parser within a NLP system 20. Alternatively still, instructions to implement the exemplary parser may be on a computer readable medium.

Introduction

The exemplary parser of a NLP system employs a goodness function to rank the possible grammatically correct parses of an utterance. The goodness function of the exemplary parser is highly accurate in representing the intended meaning of a speaker. It also has reasonable training data requirements.

With this exemplary parser, the goodness measure of a particular parse is the probability of taking each transition (“transition probability”) within the parse tree of that parse. Each transition probability within the tree is conditioned on highly predicative linguistic phenomena. Such phenomena include headwords, “phrase levels”, and “syntactic history”.

Herein, the term “linguistic features” is used to generically describe transitions, headwords (including modifying headwords), phrase levels, and syntactic history.

Statistical Goodness Measure

The statistical goodness measure (SGM) of the exemplary parser uses a generative grammar approach. In a generative grammar approach, each sentence has a top-down derivation consisting of a sequence of rule applications (i.e., transitions). The probability of the parse tree is the product of the probabilities of all the nodes. The probability for a given node is the probability that from the node one would take a specific transition, given the syntactic features.

The SGM of the exemplary parser may be calculated using either of the following equivalent formulas:

$$\text{Prob}(\text{parse}) = \prod_x \text{Prob}(\text{trn}(n_x), \text{hw}(n_y), \text{pl}(n_y), \text{sh}(n_y), \text{hw}(n_z), \text{pl}(n_z), \text{sh}(n_z) \mid \text{hw}(n_x), \text{pl}(n_x), \text{sh}(n_x), \text{segtype}(n_x))$$

Formula A

OR...

$$\text{Prob}(\text{parse}) = \prod_x \text{Prob}(\text{trn}(n_x) \mid \text{hw}(n_x), \text{pl}(n_x), \text{sh}(n_x), \text{segtype}(n_x)) \\ \text{Prob}(\text{modhw}(n_x) \mid \text{trn}(n_x), \text{hw}(n_x))$$

Formula B

where

n_x : is the X^{th} node in a parse tree

n_y & n_z : are the Y^{th} and Z^{th} nodes and children of the X^{th} node

$\text{trn}(n_x)$: is the name of the transition out of n_x of the form $X \rightarrow Y Z$

$\text{hw}(n_x)$: is the headword of n_x

$\text{pl}(n_x)$: is the phrase level of n_x

$\text{sl}(n_x)$: is the syntactic history of n_x

$\text{segtype}(n_x)$: is the segtype of n_x

$\text{modhw}(n_x)$: is the modifying headword of n_x

The exemplary parser defines phrase levels and labels them. Previous conventional approaches clustered transitions by segtype. For example, transitions focused on noun phrases, transitions focused verb phrases, etc. However, within each such grouping, the rules can be further subdivided into multiple levels. These levels are called “phrase levels” herein. These phrase levels are highly predicative of whether a transition will occur.

A null transition is utilized for each phrase level to account for no change from one level to the next. The null transition enables a node to move to the next

1 level without being altered. The null transition is assigned probabilities just like
2 other transitions.

3 The exemplary parser defines each node's syntactic history. Previous
4 conventional approaches conditioned on linguistic phenomena associated with a
5 node, its parent, and/or its children. However, such approaches are overly
6 limiting. Using the exemplary parser, phenomena that are predicative but appear
7 elsewhere in the tree (other than simply a node's immediate decedents or
8 ancestors) are included in the probability calculation.

9 The probabilities of the exemplary parser are conditioned on transition
10 name, headword, phrase level, and syntactic history.

11 Since the probabilities are conditioned on the transition name in the
12 exemplary parser instead of just the structure of the rule (e.g. $VP \rightarrow NP VP$), the
13 parser may give the same structure different probabilities. In other words, there
14 may be two transitions with the same structure that have different probabilities
15 because their transition names are different.

16 The probabilities of the SGM of the exemplary parser are computed top
17 down. This allows for an efficient and elegant method for computing the
18 goodness function.

19 A training corpus of approximately 30,000 sentences is used to initially
20 calculate the conditioned probabilities of factors such as transition name,
21 headword, syntactic bigrams, phrase level, and syntactic history. The sentences in
22 this training corpus have been annotated with ideal parse trees and the annotations
23 contain all the linguistic phenomena on which the parser conditions.
24
25

1 The probabilities computation method has two phases: training and run-
2 time. During the training phase, the system examines the training corpus, and pre-
3 computes the probabilities (which may be represented as a "count") required at
4 run-time. At run-time, the goodness function is quickly computed using these pre-
5 computed probabilities (which may be "counts").
6

7 Conditioning on Headwords

8 Consider parse trees 90 and 92 shown in Fig 5a. Assume the two parse
9 trees are identical except for the transition that created the top-most VP (verb
10 phrase).
11

12 In Tree 90 of Fig. 5a, the verb phrase was created using the rule:

13 VPwNPr1: $VP \rightarrow VP\ NP$
14

15 VPwNPr1 is used to add an object to a verb. For example, "John hit the
16 ball" or "They elected the pope."
17

18 In Tree 92 of Fig. 5a, the verb phrase was created using the rule:

19 VPwAVPr: $VP \rightarrow VP\ AVP$
20

21 VPwAVPr is used when an adverbial phrase modifies a verb. For example,
22 "He jumped high" or "I ran slowly."
23

24 To determine which tree was most probable using the conventional
25 Transition Probability Approach (TPA), which is described above in the

1 background section, the number of occurrences of VPwNPr1 and VPwAVPr in the
2 corpus is counted. If VPwNPr1 occurred most often, the conventional TPA's
3 goodness function would rank Tree 90 of Fig. 5a highest.

4 This may be correct, but often it will be wrong since it will choose Tree 90
5 of Fig. 5a regardless of the linguistic context in which the rules appear. For
6 example, assume that the headword was "smiled"

7 Parse trees 94 and 96 shown in Fig 5b illustrate the same parses shown in
8 trees 90 and 92 in Fig. 5a, but the headword "smiled" is noted.

9 English-speaking humans know that Tree 94 of Fig. 5b is highly unlikely.
10 "Smiled" is intransitive and cannot take a direct object. In other words, "She
11 smiled the ball" is incorrect because someone cannot "smile" a "ball." Although, it
12 is correct to say, "She smiled the most" because the "most" is not an object of
13 "smiled." Although "the most" can act as a noun phrase in other contexts, it is an
14 adverb in this case.

15 If the headword is included into the probability calculations, the goodness
16 function is more likely to pick the correct parse. In particular, instead of just
17 counting up all occurrences of VPwNPr1 and VPwAVPr in the corpus, a count is
18 made of how often these rules appear with the headword "smiled." In doing so, it
19 likely to be discovered that there are no instances of VPwNPr1 occurring with the
20 headword "smiled." Thus, the goodness function would calculate the probability
21 of Tree 94 of Fig. 5b to be zero.

Phrase Level

Phrases (e.g., noun phrases or verb phrases) have a natural structure. The job of the grammar (i.e., grammar rules) is to build this structure. Because of the rules of the language and because of conventions used by the grammarian, there are constraints on how the phrasal structure can be built. This translates into constraints on the order in which the rules can be applied. In other words, some rules must run before other rules. The SGM of the exemplary parser implements phrase levels to make this set of constraints explicit.

Since phrase levels are predicative of what transition can occur at each node in a parse tree, incorporating them into the goodness function makes the goodness function more accurate.

Phrase Level Defined. To define the phrase levels for a given segtype, rules that create the given segtype are grouped into levels. All the rules at a given level modify the segtype in the same way (e.g., add modifiers to the left). The levels are numbered from one to N. Each level contains a null transition that allows a node to move to the next level without having an effect on the phrase being built.

The analysis grammar build a phrase up by first producing an HW ϕ from a word. This is the head word of the phrase. It then enforces an order of levels by attaching modifiers of the headword in increasing phrase level order.

For example, consider simple noun phrases in English. When building the parse tree for a noun phrase, the determiner (e.g., “the”) is attached after the adjectives describing the noun. For example, “the red book” is correct, but “red the book” is not correct. Therefore, a rule that adds a determiner to a noun phrase

1 must come after the rule(s) that add adjectives. Again, "after" is relevant to
2 creation of a parse tree and the ordering of the application of the grammar rules.
3 The term does not relate to the order of standard writing or reading.

4 For more complex noun phrases, the grammarian building a set of rules has
5 some options. For example, consider the phrase: "The red toy with the loud
6 siren." In one set of grammar rules, the structure may be like this:

7
8 (The (red (toy (with the loud siren))))
9

10 All prepositional phrases (e.g. "with the loud siren") are attached to noun
11 first; adjectives are attached next, and finally the determiner ("the") is added last.
12 Once a determiner is attached to a noun phrase, it is not possible to add additional
13 adjectives or prepositional phrases. Another set of grammar rules might structure
14 it this way:

15
16 ((The (red toy)) (with the loud siren))
17

18 However, as long as a grammar clearly defines the structure of noun
19 phrases, there exist constraints on the order of the rules. In the exemplary parser's
20 SGM, this ordering is made explicit by adding phrase level information to the
21 rules and conditioning our probabilities on these phrase levels.

22 As another example, consider the following grammar that builds verb
23 phrases. This grammar supports verbs, noun phrases, and adjective phrases, but it
24 has been simplified and does not support a range of other valid linguistic
25

phenomena like adverbs, infinitive clauses, prepositional phrases, and conjunctions.

Rule Name	Rule	Description (and an example)
VERBtoVP	VP → VERB	Promote a verb into a verb phrase
PredAdj	VP → VP AJP	Add on an adjective. <i>"It has been <u>found effective</u>"</i>
VPwNPr1	VP → VP NP	Add direct object to a VP <i>"I <u>hit Bill</u>"</i>
Perfect	VP → VP VP	Adds the VP "have" to a VP <i>"Joe <u>had gone</u>"</i>
SubjPQuant	VP → NP VP	Add quantifiers to the VP like "all", "both", ... <i>"The children have <u>all gone</u>"</i>
VPwNPI	VP → NP VP	Add noun subject to a VP <i>"<u>John jumped.</u>"</i>
SubjectAJP	VP → AJP VP	Add AJP subject to a VP <i>"<u>More surprising is his attitude</u>"</i>
InvertAJPwS	VP → AJP VP	Add final modifier, AJP, to the left of a VP <i>"<u>More remarkable still, he went</u>"</i>
Topicalization	VP → NP VP	Add object of VP, NP, to the left of a VP <i>"<u>A better-looking man I have not seen</u>"</i>

Table 3

This grammar can parse simple verb phrases like those shown in the description column above and complex phrases like:

1
2 "More surprising, we have all found useful the guidelines which were published last year"

3
4 Fig. 6a shows a parse tree 100 representing a parse of the above sentence,
5 where the parse is done in accordance with the example grammar provided above.

6 To build complex verb phrases, this grammar enforces an ordering on the
7 rules. First, VerbtoVP always runs to create the initial verb phrase. Then, post
8 modifiers are added using PredAdj and/or VPwNPr1. Then "have" and quantifiers
9 can be added. Next, the subject is added using SubjAJP or VPwNP1. Finally,
10 topicalization and inverted AJP can be applied to phrases that have a subject

11 Constraints, such as the following, are made explicit by adding the phrase
12 level (of the exemplary parser) into the grammar rules:

Phrase Level	Rule Name	Rule	Level Description
1	VERBtoVP	VP(1) → VERB(PL_Verb_Max)	Create a VP
2	PredAdj	VP(2) → VP(1) AJP(PL_AJP_Max)	Add post modifiers
	VPwNPr1	VP(2) → VP(1) NP(PL_AJP_Max)	
	VPNull2	VP(2) → VP(1)	
3	Perfect	VP(3) → VP(1) VP(2,3)	Add "have" and quantifiers
	SubjPQuant	VP(3) → NP(PL_NP_Max) VP(2,3)	
	VPNull3	VP(3) → VP(2)	
4	VPwNPI	VP(4) → NP(PL_NP_Max) VP(3)	Add subject
		AJP(PL_AJP_Ma	
	SubjectAJP	VP(4) → x) VP(3)	
	VPNull4	VP(4) → VP(3)	
5	InvertAJPwS	VP(5) → AJP(PL_AJP_Ma x) VP(4)	Add modifiers to VPs that have a subject
	Topicalization	VP(5) → NP(PL_NP_Max) VP(4)	
	VPNull5	VP(5) → VP(4)	

Table 4

As shown above in Table 4 on the right-hand side of each rule, each constituent is associated with a particular phrase level that is required for that constituent. Specifically, the number in parenthesis indicates the phrase level of the constituent (e.g., "VP(4)").

1 On the left-hand side of the rule, the phrase level of the resulting node is
2 specified. For example, consider the null transition:

3
4
$$\text{VP}(4) \rightarrow \text{VP}(3)$$

5

6 This null transition can be applied to a VP at phrase level three and create a
7 VP at phrase level four.

8 "PL_Max" in a phrase level indicator means the highest phrase level that
9 occurs for a given segtype. For example, for the grammar above $\text{VP}(\text{PL_Max})$
10 would be the same as $\text{VP}(5)$. As another example:

11
12
$$\text{VPwNP}! : \text{VP}(4) \rightarrow \text{NP}(\text{PL_Max}) \text{VP}(3)$$

13

14 This means that the rule can be applied to an NP that is at the highest NP
15 level and to a VP that is at level three. The result of running the rule is to create a
16 VP at level four.

17 Sometimes the phrase level of a constituent of the saem segtype is the
18 resulting node and may be either at the phrase level of the resulting node or less
19 than the phrase level of the resulting node. For example:

20
21
$$\text{Perfect: VP}(3) \rightarrow \text{VP}(1) \text{VP}(2,3)$$

22

23 He melted.

24 He had melted.

25 He had been melted.

1
2
3
4 To see an example of null transitions, consider the phrase:
5

6 "Surprising, we found useful the guidelines."
7

8 Notice that this phrase differs from the similar phrase used above in that
9 "...we have all found useful..." has been simplified to be "... we found useful..."

10 The rule VpwNul at transition null requires the second constituent to have
11 PL3. Because the constituent has PL2 we construct a null transition first. Trans
12 114 to tras'.

13 Fig. 6b shows a parse tree 110 representing a parse of this sentence. The
14 null transition at 112 is used to move the VP(2) to be a VP(3). The null transition
15 can be explicitly represented in the parse tree (as shown in Fig. 6b) or be implicit.
16 It doesn't matter as long as it is taken into account in the computation of the
17 probabilities of the exemplary parser.

18 Conditioning on phrases levels means that any parse tree that violates the
19 phrase level constraints can be eliminated (given probability equal to zero) by the
20 exemplary parser.
21

22 Modeling the Syntactic Modification of Individual Words. In the
23 exemplary parser, the use of phrase levels and null transitions accurately model
24
25

1 the syntactic way that words "want" to be modified. This has not been done with
2 any conventional parser and goodness function.

3 The Penn Tree Bank is the focus of most of those working on probabilistic
4 grammars. Penn Tree Bank is annotated with parts of speech for the words and
5 minimal phrase names and brackets. The syntactic relationships between
6 constituents are not given. Without data to model, one does not develop a realistic
7 model. There is no explicit grammar given that would produce the bracketing for
8 the parses in the Penn Tree Bank. The great majority of those working on the Penn
9 Tree Bank computationally induce a grammar from the annotation. The number of
10 transitions so induced generally run to the thousands. This induction precludes null
11 transitions for two reasons.

- 12 • There is no clear hierarchy to the supplied bracketing annotations.
13 Because of this, there is no obvious way to describe a hierarchy by
14 defining in what order modifiers are to be attached to the head.
- 15 • Since a null transition could be associated with any real transition, it
16 becomes intractable, and not productive, to put all in. In addition, it's
17 not clear what subset could be induced since criteria seem to be lacking.

18 In contrast, the phrase levels and null transitions of the exemplary parser
19 models the grammar of the English natural language. For example, consider the
20 noun "nut." You would never see a sentence such as 'I want nut.' or 'Nut is on the
21 table.' The word "nut" wants a determiner such as "a" or "the". The phrase levels
22 and null transitions force the exemplary parser to explicitly consider the absence
23 of modifiers, as well as their presence.
24
25

1 Since any transition has a probability of 1 or lower, the more transitions in
2 a sentence or phrase implies a smaller goodness measure. Therefore, when
3 calculating the goodness measure using conventional approaches, the sentence "I
4 want nut" would be preferred over the sentence "I want a nut." This is because the
5 latter has more transition; therefore, the goodness measure would be less than the
6 former.

7 Since the exemplary parser considers the lack of modifiers (such as pre-
8 determiners, determiners, attributive adjectives and adverbs, post-modifiers, etc.)
9 when calculating the goodness measure, the sentence 'I want a nut' has a greater
10 goodness measure than 'I want nut.' Thus, "I want a nut" is preferred over "I want
11 nut." Although "I want a nut" appears, at a surface level, to have more transitions
12 and thus should have a lower goodness measure using conventional approaches.
13 However, using the exemplary parser the contextually correct sentence "I want a
14 nut" is preferred over "I want nut." No conventional approach has this property.
15 Another example set is 'I admonish the children.' and 'I admonish.'

16 In the exemplary parser, the transition probabilities are conditioned on
17 headwords. Using a training corpus, the exemplary parser counts up the number
18 of times a specific headword is modified by a rule and the number of times it isn't.

19 20 **Syntactic History**

21 A node's syntactic history is the relevant grammatical environment that a
22 node finds itself in. It may include the history of transitions that occur above the
23 node. For example, is the node below a NREL, PRPRT, PTPRT, RELCL, or
24 AVPVP? It may include whether the node is in a passive or an active
25

1 construction. It may include information that appears elsewhere in the tree. For
2 example, whether the headword of a sibling node is singular or plural. The
3 specifics of what it relevant is dependent upon the specifics of the grammar (i.e.,
4 rewrite rules or transitions) being used.

5 For example, Fig. 7 shows two parse trees, 120 and 130, for the same verb
6 phrase. Both trees are parsing a verb phrase having the mono-transitive headword
7 (hw="hit") and the verb phrase is known to be passive (sh=passive). In tree 120,
8 the verb has a direct object as represented by NP at 122. In tree 130, the verb does
9 not take a direct object.

10 In English, a mono-transitive verb inside a passive construction does not
11 take a direct object. In contrast, when in the active form, the mono-transitive verb
12 "hit" takes a direct object. For example, "I hit the ball" in the active form has a
13 direct object "ball" to the verb "hit", but "the ball was hit" in the passive form has
14 no direct object to "hit."

15 English-speaking humans know that tree 120 will never occur. In other
16 words, there is a zero probability of a mono-transitive verb (like "hit") taking a
17 direct object when the sentence is passive.

18 In the exemplary parser, the transition probabilities are conditioned on
19 syntactic history as well as headwords. Using a training corpus, the exemplary
20 parser counts up how often VPwNP_{r1} occurs in a passive construction with a
21 mono-transitive verb and finds that it never occurs. Thus, the probability of Tree
22 120 would be calculated to be zero.

1 Propagated Syntactic History. Syntactic history can be propagated down
2 many levels of the tree. Take, for example, the sample sentence, “Graceland, I
3 love to visit.” The thing (“Graceland”) that “I” love to visit is stated before it is
4 revealed the “I” loves to visit anything.

5 Fig. 8 shows an annotated parse tree 140 of a parse of this sample sentence.
6 As can be seen in Fig. 8, the “topicalization” feature is propagated past the verb
7 “like” to the verb “visit.”

8
9 Examples of Syntactic Phenomena. The following is a list of syntactic
10 phenomena that are incorporated in a syntactic history. This list is intended to
11 provide examples of syntactic phenomena tracked as syntactic history by the
12 exemplary parser. This list is not exclusive of other possible phenomena and is
13 not intended to be limiting.

14 These phenomena are well known by linguists, but have not been used by
15 computational linguists when considering conditional probabilities for parsers. For
16 each phenomenon, sample sentences are provided and a general description (or
17 examples) of why such phenomenon affects a phrase or sentence.

18
19 Passive.

20 Sample sentences:

- 21 • The cabbage was licked by the rabbit.
- 22 • The bomb must have been detonated from a distance.
- 23 • He was born on a log table.

24 Affect:

25

- The verb in the passive construction (e.g., “lick”, “detonate”, “bore”) in the great majority of cases does not take a syntactic object as a post-modifier.
- Passive is at the same phrase level as progressive, perfect, and modal. However, if a passive construction is being built, those are not allowed as transition below the passive and the verb that is passivized. (E.g., The bomb must be had detonated.) However, the converse is not true, passive must follow progressive, perfect, or modal.

Negative Polarity.

Sample sentences:

- Never had I seen such chaos.
- Seldom would he be home before 4 a.m.
- Rarely did he pass exams.

Affect:

- Compare the first sentence to the kernel representation: “I had never seen such chaos.” When the negation (“never”) is fronted for focusing purposes there must be subject-verb inversion: note the ‘had’ before the ‘I’.

Domodal Fronting.

Sample sentences:

- Had I seen such chaos?
- Would he be home before 4 a.m?

- Did he pass exams?

Affect:

- A question can be formed by inverting the subject and an ‘auxiliary’ verb. Then there is a lack of number agreement between the subject and the verb that follows the subject. By conditioning on DOMODAL_FRONTING the exemplary parser knows to expect this disagreement.
- There are restrictions on what can be between the fronted verb and the subject. One would not normally say “Did frequently he pass exams?” However, one could say “Frequently he passed exams.”

Comparative.

Sample sentences: (These are a small subset of the types of comparative constructions in English.)

- The artichoke has *more* brains than the rock does.
- The limpet is *more* beautiful than any defoliant can be.
- She worked *harder* than Filbert thought was possible.

Affect:

The samples above show a clause (denoted by underscore) modifying a noun, adjective, and adverb respectively. This is generally a rare construction, however it is common for constructing comparatives. For a comparative structure like this, all other post-modifiers are reduced in probability. By conditioning on the fact the grammar built a comparative structure, the exemplary parser’s can determine how well the construction follows type.

1
2 Imperative.

3 Sample sentences:

- 4 • Go to your room now.
5 • Pour three ounces of the gin into the vermouth.
6 • Please pass the trisodium phosphate.

7 Affect:

8 Sentences usually have subjects. Imperatives don't. However, it is less
9 likely that 'tensed' clauses within sentences lack subjects. By conditioning on
10 whether the parser is building a sentence or embedded clause, the parser can
11 apportion the lack of subject correctly.
12

13 Topicalization Of Verb Object. Sample sentences:

- 14 • Graceland, I love to visit.
15 • This book I must read.

16 Affect:

17 In both 'Graceland, I love to visit.' and 'I love to visit Graceland.' the verb
18 "visit" has the syntactic object "Graceland." However, in this case, a different
19 rule, at a different phrase level (*Topicalization*), is used to attach the object to the
20 verb than the usual *VPwObject*. If *Topicalization* is used then the probability of
21 *VpwObject* must be lowered when the parser get to the phrase level it operates on.
22

23 SGM of the Exemplary Parser
24
25

1 The SGM (statistical goodness measure) of the exemplary parser uses a
2 generative grammar approach—each sentence has a top-down derivation consisting
3 of a sequence of rule applications (transitions). The probability of a parse tree is
4 the product of the probabilities of all the nodes within that tree.

5 Generally, the probability of a node is defined as a conditional probability:

$$\text{Prob}(\text{node}) = \text{Prob}(\text{what_is_unknown} \mid \text{what_is_known})$$

8 Formula 1

9
10 Assume that each node is visited in a depth-first tree walk. What is known
11 is the information associated with the node and/or with any node previously
12 encountered in the tree walk. For example, the properties of the node, it is
13 headword, phrase level, syntactic history, and segtype. What is unknown is what
14 occurs below the node (i.e., the transition taken and the properties of its children).

15 Fig. 9 shows a portion of a parse tree 150 and visually illustrates what is
16 known and unknown at a node 152. What is known is above line 154 because it
17 has already been processed. Below line 154 is what is unknown because it has not
18 been processed.

19 With reference to the parse tree 150 of Fig. 9, the conditional probability of
20 exemplary parser is:

21
22 $\text{Prob}(\text{parse})$

$$23 = \prod_x \text{Prob}(n_x)$$

$$24 = \prod_x \text{Prob}(\text{trn}(n_x), \text{hw}(n_y), \text{pl}(n_y), \text{sh}(n_y), \text{hw}(n_z), \text{pl}(n_z), \text{sh}(n_z) \mid \text{hw}(n_x), \text{pl}(n_x), \text{sh}(n_x), \text{segtype}(n_x))$$

Formula 2

where n_X ranges over all nodes in the tree and the transition named by $\text{trn}(n_X)$ is of the form $X \rightarrow Y Z$ or of the form $X \rightarrow Y$.

To simplify Formula 2, it is noted that not all the parameters are independent. In particular, $\text{trn}(n_X)$ and $\text{pl}(n_X)$ imply $\text{pl}(n_Y)$ and $\text{pl}(n_Z)$. In other words, the name of the transition and the phrase level at node X implies the phrase levels of nodes Y and Z . Therefore, $\text{pl}(n_Y)$ and $\text{pl}(n_Z)$ may be removed from the left-hand side of the formula:

$$= \prod_X \text{Prob}(\text{trn}(n_X), \text{hw}(n_Y), \text{sh}(n_Y), \text{hw}(n_Z), \text{sh}(n_Z) \mid \text{hw}(n_X), \text{pl}(n_X), \text{sh}(n_X), \text{segtype}(n_X))$$

Formula 3

Similarly, Formula 3 may be simplified because $\text{trn}(n_X)$, $\text{hw}(n_X)$, and $\text{sh}(n_X)$ imply $\text{sh}(n_Y)$ and $\text{sh}(n_Z)$. In other words, the name of the transition, the headword, and the syntactic history at node X implies the syntactic history of nodes Y and Z . Therefore, $\text{sh}(n_Y)$ and $\text{sh}(n_Z)$ may be removed from the from the left-hand side of the formula:

$$= \prod_X \text{Prob}(\text{trn}(n_X), \text{hw}(n_Y), \text{hw}(n_Z) \mid \text{hw}(n_X), \text{pl}(n_X), \text{sh}(n_X), \text{segtype}(n_X))$$

Formula 4

Formula 4 may be further simplified. Tracking *both* $\text{hw}(n_Y)$ and $\text{hw}(n_Z)$ is not particularly valuable because one of them is the same as $\text{hw}(n_X)$. The one that

1 is not the same is the modifying headword. The notation $\text{modhw}(n_x)$ to refer to
2 this modifying headword. This yields:

$$3 \quad = \Pi_x \text{Prob}(\text{trn}(n_x), \text{modhw}(n_x) \mid \text{hw}(n_x), \text{pl}(n_x), \text{sh}(n_x), \text{segtype}(n_x))$$

4
5 Formula 5

6
7 Formula 5 may be simplified still further by applying the chain rule (as
8 understood by those skilled in the art of statistics), yields this:

$$9 \quad = \Pi_x \text{Prob}(\text{trn}(n_x) \mid \text{hw}(n_x), \text{pl}(n_x), \text{sh}(n_x), \text{segtype}(n_x)) * \\ 10 \quad \text{Prob}(\text{modhw}(n_x) \mid \text{trn}(n_x), \text{hw}(n_x), \text{pl}(n_x), \text{sh}(n_x), \text{segtype}(n_x))$$

11
12 Formula 6

13
14 Since $\text{trn}(n_x)$ implies $\text{pl}(n_x)$ and $\text{segtype}(n_x)$, the Formula 6 can further
15 simplify this to be:

$$16 \quad = \Pi_x \text{Prob}(\text{trn}(n_x) \mid \text{hw}(n_x), \text{pl}(n_x), \text{sh}(n_x), \text{segtype}(n_x)) * \\ 17 \quad \text{Prob}(\text{modhw}(n_x) \mid \text{trn}(n_x), \text{hw}(n_x), \text{sh}(n_x))$$

18
19 Formula 7

20
21 Finally, since it has been found that $\text{sh}(n_x)$ is not very predicative of what
22 the modifying headword will be, Formula 7 can be approximated by removing
23 $\text{sh}(n_x)$ from that part of Formula 7:

$$\cong \prod_X \text{Prob}(\text{trn}(n_X) \mid \text{hw}(n_X), \text{pl}(n_X), \text{sh}(n_X), \text{segtype}(n_X)) \text{Prob}(\text{modhw}(n_X) \mid \text{trn}(n_X), \text{hw}(n_X))$$

Formula 8 (SGM for a parse)

Notice that Formula 8 above is Formula B recited near the beginning of this detailed description.

PredParamRule Probability and SynBigram Probability

As described above, the probability of a parse tree is the products of the probabilities of each node. The probability of each node is the product of two probabilities. Thus, the SGM probability formula for a single node in a tree may be rewritten like this:

$$\text{Prob}(\text{trn}(n_X) \mid \text{hw}(n_X), \text{pl}(n_X), \text{sh}(n_X), \text{segtype}(n_X)) \text{Prob}(\text{modhw}(n_X) \mid \text{trn}(n_X), \text{hw}(n_X))$$

Formula 9 (SGM probability at a given node X)

where X ranges over all the nodes in the parse tree.

This represents the statistical goodness measure (SGM) of the exemplary parser. This may be divided into to two parts. For convenience, the first probability will be called the predictive-parameter-and-rule probability or simply “PredParamRule Probability” and the second probability will be called the “SynBigram Probability”.

The PredParamRule Probability is:

$$\text{Prob}(\text{trn}(n_X) \mid \text{hw}(n_X), \text{pl}(n_X), \text{sh}(n_X), \text{segtype}(n_X))$$

Formula 10 (PredParamRule Probability)

1 Unlike the Simple Content Dependent Approach (described above in the
2 background section), the PredParamRule Probability of the exemplary parser
3 conditions upon headword, segtype, phrase level, and syntactic history. Since
4 these are highly predicative of the contextually correct parse, this PredParamRule
5 Probability is a significantly more accurate goodness function than conventional
6 techniques.

7 The SynBigram Probability is:

$$\text{Prob}(\text{modhw}(n_x) \mid \text{trn}(n_x), \text{hw}(n_x))$$

8
9
10 **Formula 11 (SynBigram Probability)**

11
12 The SynBigram Probability computes the probability of a syntactic bigram.
13 Syntactic bigrams are two-word collocation. The probability a measure of the
14 “strength” of the likelihood of a pair of words appearing together in a syntactic
15 relationship. For example, the object of the verb “drink” is more likely to be
16 “coffee” or “water” than “house”.

17 As described above in the background section, this is a conventional
18 technique to calculate a goodness measure. However, with existing conventional
19 syntactic bigram approaches, it is used alone to calculate the goodness function
20 and it requires a huge training corpus.

21 The exemplary parser overcomes the limitations of conventional syntactic
22 bigram approaches by further conditioning the goodness measure on independent
23 probability characteristics. In particular, those characteristics are represented by
24 the PredParamRule Probability formula (Formula 10).
25

As a review, the following is a known about calculating conditional probabilities by counting appearances in a training corpus:

$$\begin{aligned} \text{Prob}(x|y) &= \frac{\text{Prob}(x\&y)}{\text{Prob}(y)} \\ &= \frac{\text{Count}(x\&y)}{\text{Count}(y)} \end{aligned}$$

Therefore, the PredParamRule Probability and the SynBigram Probability can be calculated by counting the appearances of relevant events in the training corpus. The probabilities of a given training corpus that are determined by the PredParamRule Probability and the SynBigram Probability may be generally called "language-usage probabilities" for that given training corpus.

Thus, the PredParamRule Probability formula (Formula 10) may be calculated as follows:

PredParamRule Probability

$$\begin{aligned} &= \text{Prob}(\text{trn}(n_x) \mid \text{hw}(n_x), \text{pl}(n_x), \text{sh}(n_x), \text{segtype}(n_x)) \\ &= \frac{\text{Count}(\text{trn}(n_x) \& \text{hw}(n_x) \& \text{pl}(n_x) \& \text{sh}(n_x) \& \text{segtype}(n_x))}{\text{Count}(\text{hw}(n_x) \& \text{pl}(n_x) \& \text{sh}(n_x) \& \text{segtype}(n_x))} \end{aligned}$$

Formula 12

Moreover, the SynBigram Probability formula (Formula 11) may be calculated as follows:

SynBigram Probability

$$\begin{aligned} &= \text{Prob}(\text{modhw}(n_x) \mid \text{trn}(n_x), \text{hw}(n_x)) \\ &= \frac{\text{Count}(\text{modhw}(n_x) \ \& \ \text{trn}(n_x) \ \& \ \text{hw}(n_x))}{\text{Count}(\text{trn}(n_x) \ \& \ \text{hw}(n_x))} \end{aligned}$$

Formula 13

Two Phases of SGM Calculation

Typically, a parser of an NLP system (such as the exemplary parser) is designed to quickly calculate the goodness measure for many parse trees of parses of a phrase. To accomplish this, the exemplary parser is implemented in two phases: “training” and “run-time.”

During the training phase, the exemplary parser pre-calculates the counts that are needed to compute the PredParamRule Probability and the SynBigram Probability at run-time. Although this process tends to be time-consuming, processor-intensive, and resource-intensive, it only need be once for a given training corpus.

The result of the training phase is a set of counts for headword, phrase level, syntactic history, and segtype. If the training corpus approximates the

1 natural language usage of a given purpose (general, specific, or customized), then
2 the counts also approximate the natural language usage for the same purpose.

3 At run-time, these pre-calculated counts are used to quickly determine the
4 probability of the parse tree. Each phrase is parsed into multiple parse trees. Each
5 parse tree is given a SGM based upon the pre-calculated counts.

6 Alternatively, the training and run-time phase may be performed nearly
7 concurrently. The training phase may be performed on a training corpus (or some
8 subset of such corpus) just before the run-time phase is performed. Those who are
9 skilled in the art will understand that time and space trade-offs may be made to
10 accommodate the given situation.

11 Regardless, the training phase (or some portion thereof) is performed, at
12 least momentarily, before the run-time phase. This is because the training phase
13 provides the foundation for the run-time phase to base its SGM calculations.

14 **Training Phase**

15
16 Fig. 10 shows a methodological implementation of the training phase of the
17 exemplary parser. The training phase has two parts: the preparation part and the
18 computation part. The preparation part is performed before the computation part.

19 During the preparation part, a training corpus is created at 200. The
20 training corpus includes a body of "correctly" parsed sentences (and phrases) that
21 the parser can use to determine correct goodness measures for similarly structured
22 sentences. At 202, the parser examines each parse tree of the "correctly" parsed
23 sentences (and phrases) of the corpus. To examine the corpus, a depth-first tree
24
25

1 walk is performed on each parse tree. At 204, the syntactic history for each node
2 is computed and stored in the node.

3 After the training corpus is created in the preparation part, the computation
4 part of the training phase begins at 206. This part computes all the counts, at 206,
5 used in the PredParamRule Probability and SynBigram Probability. In other
6 words, the exemplary parser counts how often each of the following combinations
7 are seen in the training corpus:

8
9 (transition, headword, phrase level, syntactic history, segtype)

10 (headword, phrase level, syntactic history, segtype)

11 (modifying headword, transition, headword)

12 (transition, headword)

13
14 In the exemplary implementation of the parser, this is done by creating four
15 multi-dimensional arrays -- one for each set of counts. In particular, the following
16 arrays are used:

- 17
18 • RuleCountNumerator: Each entry stores the count for a different
19 combination of the transition, headword, phrase level, syntactic
20 history, and segtype. These counts are used in the numerator of the
21 PredParamRule Probability.
22 • RuleCountDenominator: Each entry stores the count for a different
23 combination of headword, phrase level, syntactic history, and
24
25

segtype. These counts are used in the denominator of the PredParamRule Probability.

- BigramCountNumerator: Each entry stores the count for a different combination of modifying headword, transition, and headword. These counts are used in the numerator of the SynBigram Probability.
- BigramCountDenominator: Each entry stores the count for a different combination of transition and headword. These counts are used in the denominator of the SynBigram Probability.

In the exemplary implementation of the parser, the arrays initially contain all zeros. All nodes in the training corpus are examined and the corresponding entries in the arrays are incremented. At 208, the results are stored. At 210, the process ends.

This process may be described by the following exemplary pseudocode:

For each parse tree t in the training corpus

For each node n in t

RuleCountNumerator(trn(n), hw(n), pl(n), sh(n), segtype(n)) = +1;

RuleCountDenominator(hw(n), pl(n), sh(n), segtype(n)) = +1;

BigramCountNumerator(modhw(n), trn(n), hw(n)) = +1;

BigramCountDenominator(trn(n), hw(n)) = +1;

End loop

End loop

1 If there is no modifying headword, as happens with unary rules, then there
2 is no counting performed.

3 4 Run-time Phase

5 Given the counts computed in the training phase (described above), the
6 goodness measure for a given parse tree may be calculated quickly and efficiently.
7 The probability of the tree is the product of the probabilities of the nodes and the
8 probability of each node is quickly computed using the pre-computed counts.

9 Fig. 11 shows a methodological implementation of the run-time phase of
10 the exemplary parser. At 300, a parse of a phrase is initiated. An application,
11 such as a grammar checker, may initiate such a parse by the exemplary parser. At
12 302, the exemplary parser parses the phrase and generates one or more parse trees.
13 Each tree represents a grammatically valid parse of the phrase. If there is only one
14 valid parse, there is no need to rank it for apparent reasons. Thus, this process
15 may jump ahead to blocks 316 to report the results and 318 to end the process.

16 At 304, the exemplary parser examines and calculates the SGM for the first
17 of the parse trees. The order in which the trees are examined does not affect the
18 results. Therefore, any tree in the set of valid parse trees may be the first. Blocks
19 306-312 show the details of examinations and SGM calculations for all of the trees
20 of a phrase.

21 At 306, the exemplary parser calculates the probability at each node in the
22 by using PredParamRule Probability and SynBigram Probability. To do this, the
23 exemplary parser uses the counts from the training phase (described above and
24 shown in Fig. 10).
25

At 308 in Fig. 11, the exemplary parser calculates probability (i.e., SGM) of the tree as a product of the probabilities of the nodes in the tree. At 310, the process determines if there are more trees to be examined. If so, then the process examines the next tree at 312 and then loops back through blocks 306-310. If all trees of a phrase have been examined, then the exemplary parser ranks each parse tree based upon their SGM at 314.

At 316, the exemplary parser does something with the results. It may store it, report it, return it, display it, or the like. The run-time process ends at 318.

This process may be described by the following exemplary pseudocode:

goodness=1;

For each node n in the parse tree

$$\text{rule_prob} = \frac{\text{RuleCountNumerator}(\text{trn}(n), \text{hw}(n), \text{pl}(n), \text{sh}(n), \text{segtype}(n))}{\text{RuleCountDenominator}(\text{hw}(n), \text{pl}(n), \text{sh}(n), \text{segtype}(n))}$$

$$\text{bigram_prob} = \frac{\text{BigramCountNumerator}(\text{modhw}(n), \text{trn}(n), \text{hw}(n))}{\text{BigramCountDenominator}(\text{trn}(n), \text{hw}(n))}$$

goodness = goodness * rule_prob * bigram_prob;

End loop

Alternatives

Computing More Precise Bigrams. Above, Formula 7 was simplified into Formula 8 by removing the syntactic history of n_X (i.e., $\text{sh}(n_X)$). This is done because $\text{sh}(n_X)$ is not very predicative of what the modifying headword will be.

1 Thus, Formula 7 may be approximated by removing $sh(n_X)$. The result is Formula
2 8.

3 This is a reasonable simplification. However, it should be clear to anyone
4 of ordinary skill in the art that this simplification is not necessary. If one keeps
5 $sh(n_X)$ in the formula (as is the case in Formula 7), then the probabilities will be
6 more accurate. However, a larger and more accurate training corpus is necessary.
7

8 Making the Training Phase More Efficient. As described, the training
9 phase has two parts: the preparation and the calculation. However, it should be
10 clear to anyone of ordinary skill in the art that the training phase may be
11 accomplished in one part by merging the steps of each part into a single pass over
12 the training corpus.
13

14 Word Classes. In general, syntactic bigrams requires a great deal of
15 training data. To make the statistics gathering more tractable, words could be
16 grouped into clusters with similar distributional properties.
17

18 Exemplary Computing Environment

19 Fig. 12 illustrates an example of a suitable computing environment 920 on
20 which the exemplary ranking parser may be implemented. The exemplary
21 computing environment 920 may be a computing environment comprising or
22 utilizing a NLP system.

23 Exemplary computing environment 920 is only one example of a suitable
24 computing environment and is not intended to suggest any limitation as to the
25

1 scope of use or functionality of the invention. Neither should the computing
2 environment 920 be interpreted as having any dependency or requirement relating
3 to any one or combination of components illustrated in the exemplary computing
4 environment 920.

5 The exemplary ranking parser is operational with numerous other general
6 purpose or special purpose computing system environments or configurations.
7 Examples of well known computing systems, environments, and/or configurations
8 that may be suitable for use with the exemplary ranking parser include, but are not
9 limited to, personal computers, server computers, thin clients, thick clients, hand-
10 held or laptop devices, multiprocessor systems, microprocessor-based systems, set
11 top boxes, programmable consumer electronics, network PCs, minicomputers,
12 mainframe computers, distributed computing environments that include any of the
13 above systems or devices, and the like.

14 The exemplary ranking parser may be described in the general context of
15 computer-executable instructions, such as program modules, being executed by a
16 computer. Generally, program modules include routines, programs, objects,
17 components, data structures, etc. that perform particular tasks or implement
18 particular abstract data types. The exemplary ranking parser may also be practiced
19 in distributed computing environments where tasks are performed by remote
20 processing devices that are linked through a communications network. In a
21 distributed computing environment, program modules may be located in both local
22 and remote computer storage media including memory storage devices.

23 As shown in Fig. 12, the computing environment 920 includes a general-
24 purpose computing device in the form of a computer 930. The components of
25

1 computer 920 may include, by are not limited to, one or more processors or
2 processing units 932, a system memory 934, and a bus 936 that couples various
3 system components including the system memory 934 to the processor 932.

4 Bus 936 represents one or more of any of several types of bus structures,
5 including a memory bus or memory controller, a peripheral bus, an accelerated
6 graphics port, and a processor or local bus using any of a variety of bus
7 architectures. By way of example, and not limitation, such architectures include
8 Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA)
9 bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA)
10 local bus, and Peripheral Component Interconnects (PCI) bus also known as
11 Mezzanine bus.

12 Computer 930 typically includes a variety of computer readable media.
13 Such media may be any available media that is accessible by computer 930, and it
14 includes both volatile and non-volatile media, removable and non-removable
15 media.

16 In Fig. 12, the system memory includes computer readable media in the
17 form of volatile, such as random access memory (RAM) 940, and/or non-volatile
18 memory, such as read only memory (ROM) 938. A basic input/output system
19 (BIOS) 942, containing the basic routines that help to transfer information
20 between elements within computer 930, such as during start-up, is stored in ROM
21 938. RAM 940 typically contains data and/or program modules that are
22 immediately accessible to and/or presently be operated on by processor 932.

23 Computer 930 may further include other removable/non-removable,
24 volatile/non-volatile computer storage media. By way of example only, Fig. 12
25

1 illustrates a hard disk drive 944 for reading from and writing to a non-removable,
2 non-volatile magnetic media (not shown and typically called a "hard drive"), a
3 magnetic disk drive 946 for reading from and writing to a removable, non-volatile
4 magnetic disk 948 (e.g., a "floppy disk"), and an optical disk drive 950 for reading
5 from or writing to a removable, non-volatile optical disk 952 such as a CD-ROM,
6 DVD-ROM or other optical media. The hard disk drive 944, magnetic disk drive
7 946, and optical disk drive 950 are each connected to bus 936 by one or more
8 interfaces 954.

9 The drives and their associated computer-readable media provide
10 nonvolatile storage of computer readable instructions, data structures, program
11 modules, and other data for computer 930. Although the exemplary environment
12 described herein employs a hard disk, a removable magnetic disk 948 and a
13 removable optical disk 952, it should be appreciated by those skilled in the art that
14 other types of computer readable media which can store data that is accessible by a
15 computer, such as magnetic cassettes, flash memory cards, digital video disks,
16 random access memories (RAMs), read only memories (ROM), and the like, may
17 also be used in the exemplary operating environment.

18 A number of program modules may be stored on the hard disk, magnetic
19 disk 948, optical disk 952, ROM 938, or RAM 940, including, by way of example,
20 and not limitation, an operating system 958, one or more application programs
21 960, other program modules 962, and program data 964.

22 A user may enter commands and information into computer 930 through
23 input devices such as keyboard 966 and pointing device 968 (such as a "mouse").
24 Other input devices (not shown) may include a microphone, joystick, game pad,
25

1 satellite dish, serial port, scanner, or the like. These and other input devices are
2 connected to the processing unit 932 through an user input interface 970 that is
3 coupled to bus 936, but may be connected by other interface and bus structures,
4 such as a parallel port, game port, or a universal serial bus (USB).

5 A monitor 972 or other type of display device is also connected to bus 936
6 via an interface, such as a video adapter 974. In addition to the monitor, personal
7 computers typically include other peripheral output devices (not shown), such as
8 speakers and printers, which may be connected through output peripheral interface
9 975.

10 Computer 930 may operate in a networked environment using logical
11 connections to one or more remote computers, such as a remote computer 982.
12 Remote computer 982 may include many or all of the elements and features
13 described herein relative to computer 930.

14 Logical connections shown in Fig. 12 are a local area network (LAN) 977
15 and a general wide area network (WAN) 979. Such networking environments are
16 commonplace in offices, enterprise-wide computer networks, intranets, and the
17 Internet.

18 When used in a LAN networking environment, the computer 930 is
19 connected to LAN 977 network interface or adapter 986. When used in a WAN
20 networking environment, the computer typically includes a modem 978 or other
21 means for establishing communications over the WAN 979. The modem 978,
22 which may be internal or external, may be connected to the system bus 936 via the
23 user input interface 970, or other appropriate mechanism.
24
25

1 Depicted in Fig. 12, is a specific implementation of a WAN via the Internet.
2 Over the Internet, computer 930 typically includes a modem 978 or other means
3 for establishing communications over the Internet 980. Modem 978, which may
4 be internal or external, is connected to bus 936 via interface 970.

5 In a networked environment, program modules depicted relative to the
6 personal computer 930, or portions thereof, may be stored in a remote memory
7 storage device. By way of example, and not limitation, Fig. 12 illustrates remote
8 application programs 989 as residing on a memory device of remote computer
9 982. It will be appreciated that the network connections shown and described are
10 exemplary and other means of establishing a communications link between the
11 computers may be used.
12

13 **Exemplary Operating Environment**

14 Figs. 1 and 12 illustrate examples of a suitable operating environments (20
15 in Fig. 1 and 930 in Fig. 12) in which the exemplary ranking parser may be
16 implemented. Specifically, the exemplary ranking parser is implemented by the
17 parse ranker 34 in Fig 1 and by any program 960-962 or operating system 958 in
18 Fig. 12.

19 The operating environments are only examples of suitable operating
20 environments and are not intended to suggest any limitation as to the scope of use
21 of functionality of the ranking parser described herein. Other well known
22 computing systems, environments, and/or configurations that may be suitable for
23 use with the ranking parser include, but are not limited to, personal computers,
24 server computers, hand-held or laptop devices, multiprocessor systems,
25

1 microprocessor-based systems, programmable consumer electronics, network PCs,
2 minicomputers, mainframe computers, distributed computing environments that
3 include any of the above systems or devices, and the like.
4

5 **Computer-Executable Instructions**

6 An implementation of the exemplary ranking parser may be described in
7 the general context of computer-executable instructions, such as program modules,
8 executed by one or more computers or other devices. Generally, program modules
9 include routines, programs, objects, components, data structures, etc. that perform
10 particular tasks or implement particular abstract data types. Typically, the
11 functionality of the program modules may be combined or distributed as desired in
12 various embodiments.
13

14 **Computer Readable Media**

15 An implementation of the exemplary ranking parser may be stored on or
16 transmitted across some form of computer readable media. Computer readable
17 media can be any available media that can be accessed by a computer. By way of
18 example, and not limitation, computer readable media may comprise computer
19 storage media and communications media.

20 Computer storage media include volatile and non-volatile, removable and
21 non-removable media implemented in any method or technology for storage of
22 information such as computer readable instructions, data structures, program
23 modules, or other data. Computer storage media includes, but is not limited to,
24 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
25

1 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
2 tape, magnetic disk storage or other magnetic storage devices, or any other
3 medium which can be used to store the desired information and which can be
4 accessed by a computer (particularly a computer implementing a NLP system).

5 Communication media typically embodies computer readable instructions,
6 data structures, program modules, or other data in a modulated data signal such as
7 carrier wave or other transport mechanism and included any information delivery
8 media. The term "modulated data signal" means a signal that has one or more of
9 its characteristics set or changed in such a manner as to encode information in the
10 signal. By way of example, and not limitation, communication media includes
11 wired media such as a wired network or direct-wired connection, and wireless
12 media such as acoustic, RF, infrared, and other wireless media. Combinations of
13 any of the above are also included within the scope of computer readable media.

14 **Conclusion**

15
16 Although the ranking parser for NLP has been described in language
17 specific to structural features and/or methodological steps, it is to be understood
18 that the ranking parser defined in the appended claims is not necessarily limited to
19 the specific features or steps described. Rather, the specific features and steps are
20 disclosed as preferred forms of implementing the claimed ranking parser.